arXiv:1411.3436v1 [stat.ML] 13 Nov 2014

# SelfieBoost: A Boosting Algorithm for Deep Learning

Shai Shalev-Shwartz*

**Abstract**

We describe and analyze a new boosting algorithm for deep learning called SelfieBoost. Unlike other boosting algorithms, like AdaBoost, which construct ensembles of classifiers, SelfieBoost boosts the accuracy of a single network. We prove a $\log(1/\epsilon)$ convergence rate for SelfieBoost under some "SGD success" assumption which seems to hold in practice.

## 1 Introduction

Deep learning, which involves training artificial neural networks with many layers, becomes one of the most significant recent developments in machine learning. Deep learning have shown very impressive practical performance on a variety of domains (e.g. [13, 9, 18, 2, 5, 16, 12, 11, 26, 6, 3]).

One of the most successful approaches for training neural networks is Stochastic Gradient Descent (SGD) and its variants (see for example [14, 4, 1, 15, 24, 25, 21]). The two main advantages of SGD are the constant cost of each iteration (which does not depend on the number of examples) and the ability to overcome local minima. However, a major disadvantage of SGD is its slow convergence.

There have been several attempts to speed up the convergence rate of SGD, such as momentum [24], second order information [17, 7], and variance reducing methods [23, 19, 10].

Another natural approach is to use SGD as a weak learner and apply a boosting algorithm such as AdaBoost [8]. See for example [20]. The celebrated analysis of AdaBoost guarantees that if at each boosting iteration the weak learner manages to produce a classifier which is slightly better than a random guess than after $O(\log(1/\epsilon))$ iterations, AdaBoost returns an ensemble of classifiers whose training error is at most $\epsilon$.

However, a major disadvantage of AdaBoost is that it outputs an ensemble of classifiers. In the context of deep learning, this means that at prediction time, we need to apply several neural networks on every example to make the prediction. This leads to a rather slow predictor.

In this paper we present the SelfieBoost algorithm, which boosts the performance of a single network. SelfieBoost can be applied with SGD as its weak learner, and we prove that its convergence rate is $O(\log(1/\epsilon))$ provided that SGD manages to find a constant accuracy solution at each boosting iteration.

## 2 The SelfieBoost Algorithm

We now describe the SelfieBoost algorithm. We focus on a binary classification problem in the realizable case. Let $(x_1, y_1), \ldots, (x_m, y_m)$ be the training set, with $x_i \in \mathbb{R}^d$ and $y_i \in \{\pm 1\}$. Let $\mathcal{H}$ be the class of all functions that can be implemented by a neural network of a certain architecture. We assume that there

*School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel. This research is supported by Intel (ICRI-CI).

exists a network $f^* \in \mathcal{H}$ such that $y_i f^*(x_i) \geq 1$ for all $i \in [m]$. Similarly to the AdaBoost algorithm, we maintain weights over the examples that are proportional to minus the signed margin, $y_i f_t(x_i)$. This focuses the algorithm on the hard cases. Focusing the learner on the mistakes of $f_t$ might cause $f_{t+1}$ to work well on these examples but deteriorate on the examples on which $f_t$ performs well. AdaBoost overcomes this problem by remembering all the intermediate classifiers and predicting a weighted majority of their predictions. In contrast, SelfieBoost forgets the intermediate classifiers and outputs just the last classifier. We therefore need another method to make sure that the performance does not deteriorate. SelfieBoost achieves this goal by regularizing $f_{t+1}$ so that its predictions will not be too far from the predictions of $f_t$.

**SelfieBoost**

**Parameters:** Edge parameter, $\rho \in (0, \frac{1}{4})$, and number of iterations $T$
**Initialization:** Start with an initial network $f_1$ (e.g. by running few SGD iterations)
**for** $t = 1, \ldots, T$
  define weights over the $m$ examples according to $D_i \propto e^{-y_i f_t(x_i)}$
  let $S$ be $n$ indices chosen at random according to the distribution $D$
  use SGD for approximately finding a network as follows:

$$f_{t+1} \approx \operatorname*{argmin}_g \sum_{i \in S} y_i (f_t(x_i) - g(x_i)) + \frac{1}{2} \sum_{i \in S} (g(x_i) - f_t(x_i))^2$$

  **if** $\sum_{i=1}^m D_i \left[ y_i (f_t(x_i) - f_{t+1}(x_i)) + \frac{1}{2}(g(x_i) - f_t(x_i))^2 \right] < -\rho$ and $\forall i, y_i (f_{t+1}(x_i) - f_t(x_i)) \leq 1$
    continue to next iteration
  **else**
    increase the number of SGD iterations and/or the architecture and try again
    break if no such $f_{t+1}$ found

# 3 Analysis

For any classifier, define

$$M(f) = \left| \{ i \in [m] : y_i f(x_i) \leq 0 \} \right| ,$$

to be the number of mistakes the classifier makes on the training examples and

$$\text{err}(f) = \frac{M(f)}{m}$$

to be the error rate.

Our first theorem bounds $\text{err}(f)$ using the number of SelfieBoost iterations and the edge parameter $\rho$.

**Theorem 1.** *Suppose that SelfieBoost (for either the realizable or unrealizable cases) is run for $T$ iterations with an edge parameter $\rho$. Then, the number of mistakes of $f_{T+1}$ is bounded by*

$$\textit{err}(f_{T+1}) \;\leq\; e^{-\rho^2 T} .$$

*In other words, for any $\epsilon > 0$, if SelfieBoost performs*

$$T \geq \frac{\log(1/\epsilon)}{\rho^2}$$

*iterations then we must have $\textit{err}(f_{T+1}) \leq \epsilon$.*

*Proof.* Define $\ell_i(f) = 1 - y_i f(x_i)$, $L_i(f) = \ell_i(f) - 1$ and

$$L(f) = \log\left(\sum_i \exp(L_i(f))\right) .$$

Observe that $\exp(L_i(f))$ upper bounds the zero-one loss of $f$ on $(x_i, y_i)$, and so $\log(M(f)) \leq L(f)$.

Suppose that we start with $f_1$ such that $L(f_1) \leq \log(m)$. For example, we can take $f_1 \equiv 0$. At each iteration of the algorithm, we find $f_{t+1}$ such that $\forall i, \ell_i(f_t) - \ell_i(f_{t+1}) \leq 1$ and

$$\sum_{i=1}^{m} D_i \left[\ell_i(f_t) - \ell_i(f_{t+1}) + \frac{1}{2}(\ell_i(f_t) - \ell_i(f_{t+1}))^2\right] < -\rho \, . \tag{1}$$

We will show that for such $f_{t+1}$ we have that $L(f_{t+1}) - L(f_t) < \rho$, which implies that after $T$ rounds we'll have $L(f_{T+1}) \leq L(f_1) - \rho T \leq \log(m) - \rho T$. But, we also have that $L(f_{T+1}) \geq \log(M(f_{T+1}))$, hence

$$\log(M(f_{T+1})) \leq \log(m) - \rho T \quad \Rightarrow \quad \text{err}(f_{T+1}) = \frac{M(f_{T+1})}{m} \;\leq\; e^{-\rho T} \, ,$$

as required.

It is left to show that (1) indeed upper bounds $L(f_{t+1}) - L(f_t)$. To see this, we rely on the following bound, which holds for every vectors $\theta, \lambda$ for which $\theta_i - \lambda_i \leq 1$ for all $i$ (see for example the proof of Theorem 2.22 in [22]):

$$\log(\sum_i e^{\lambda_i}) \leq \log(\sum_i e^{\theta_i}) + \sum_i \frac{e^{\theta_i}}{Z}(\lambda_i - \theta_i) + \frac{1}{2}\sum_i \frac{e^{\theta_i}}{Z}(\lambda_i - \theta_i)^2$$

Therefore, with $\lambda_i = L_i(g)$, $\theta_i = L_i(f)$, and $D_i = e^{L_i(f)}/Z$, we have

$$\begin{aligned} L(g) - L(f) &= \log(\sum_i e^{\lambda_i}) - \log(\sum_i e^{\theta_i}) \\ &\leq \sum_i D_i(L_i(g) - L_i(f)) + \frac{1}{2}\sum_i D_i(L_i(g) - L_i(f))^2 \\ &= \sum_i D_i(\ell_i(g) - \ell_i(f)) + \frac{1}{2}\sum_i D_i(\ell_i(g) - \ell_i(f))^2 \, . \end{aligned} \tag{2}$$

So, at each iteration, our algorithm will try to find $g$ that approximately minimizes the right-hand side. In the next sub-sections we show that there exists $g$, which can be implemented as a simple network, that achieves a sufficiently small objective.

□

Note that we obtain behavior which is very similar to AdaBoost, in which the number of iterations depends on an "edge" — here the "edge" is how good we manage to minimize (1).

So far, we have shown that if SelfieBoost performs enough iterations then it will produce an accurate classifier. Next, we show that it is indeed possible to find a network with a small value of (1). The key lemma below shows that SelfieBoost can progress.

**Lemma 1.** *Let $f$ be a network of size $k_1$ and $f^*$ be a network of size $k_2$. Assume that err$(f^*) = 0$. Then, there exists a network $g$ of size $k_1 + k_2 + 1$ such that $\forall i, \ell_i(f) - \ell_i(g) \leq 1$ and (2) is at most $-1/2$.*

*Proof.* Choose $g$ s.t. $g(x_i) = f(x_i) + y_i = f(x_i) + f^*(x_i)$. Clearly, the size of $g$ is $k_1 + k_2 + 1$. In addition, for every $i$,

$$\ell_i(g) - \ell_i(f) = -y_i g(x_i) + y_i f(x_i) = -y_i^2 = -1 \ .$$

This implies that (2) becomes

$$-\sum_i D_i + \frac{1}{2}\sum_i D_i = -\frac{1}{2}\sum_i D_i = -\frac{1}{2} \ .$$

□

The above lemma tells us that at each iteration of SelfieBoost, it is possible to find a network for which the objective value of (1) is at most $-1/2$. Since SGD can rather quickly find a network whose objective is bounded by a constant from the optimum, we can use, say, $\rho = 0.1$ and expect that SGD will find a network with (1) smaller than $-\rho$.

Observe that when we apply SGD we can either sample an example according to $D$ at each iteration, or sample $n$ indices according to $D$ and then let the SGD sample uniformly from these $n$ indices.

**Remark:** The above lemma shows us that we might need to increase the network by the size of $f^*$ at each SelfieBoost iterations. In practice, we usually learn networks which are significantly larger than $f^*$ (because this makes the optimization problem SGD solves easier). Therefore, one may expect that even without increasing the size of the network we'll be able to find a new network with (1) being negative.

# 4 Discussion and Future Work

We have described ana analyzed a new boosting algorithm for deep learning, whose main advantage is the ability to boost the performance of the same network.

In the future we plan to evaluate the empirical performance of SelfieBoost for challenging deep learning tasks. We also plan to generalize the results beyond binary classification problems.

# References


[1] Y. Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1): 1–127, 2009.

[2] Y. Bengio and Y. LeCun. Scaling learning algorithms towards ai. *Large-Scale Kernel Machines*, 34, 2007.

[3] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:1798–1828, 2013.

[4] Olivier Bousquet and Léon Bottou. The tradeoffs of large scale learning. In *Advances in neural information processing systems*, pages 161–168, 2008.

[5] R. Collobert and J. Weston. A unified architecture for natural language processing: deep neural networks with multitask learning. In *ICML*, 2008.

[6] G. Dahl, T. Sainath, and G. Hinton. Improving deep neural networks for lvcsr using rectified linear units and dropout. In *ICASSP*, 2013.

[7] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.

[8] Y. Freund and R.E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *European Conference on Computational Learning Theory (EuroCOLT)*, pages 23–37. Springer-Verlag, 1995.

[9] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006.

[10] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.

[11] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[12] Q. V. Le, M.-A. Ranzato, R. Monga, M. Devin, G. Corrado, K. Chen, J. Dean, and A. Y. Ng. Building high-level features using large scale unsupervised learning. In *ICML*, 2012.

[13] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361, 1995.

[14] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[15] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.

[16] H. Lee, R. Grosse, R. Ranganath, and A.Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML*, 2009.

[17] James Martens. Deep learning via hessian-free optimization. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 735–742, 2010.

[18] M.A. Ranzato, F.J. Huang, Y.L. Boureau, and Y. Lecun. Unsupervised learning of invariant feature hierarchies with applications to object recognition. In *CVPR*, 2007.

[19] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *arXiv preprint arXiv:1309.2388*, 2013.

[20] Holger Schwenk and Yoshua Bengio. Boosting neural networks. *Neural Computation*, 12(8):1869–1887, 2000.

[21] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

[22] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

[23] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.

[24] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013.

[25] Ilya Sutskever. *Training recurrent neural networks*. PhD thesis, University of Toronto, 2013.

[26] M. Zeiler and R. Fergus. Visualizing and understanding convolutional neural networks. *arXiv preprint arXiv:1311.2901*, 2013.